



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/698,265	10/30/2003	Magnus Karlsson	200312448-1	2953
22879	7590	01/04/2006	EXAMINER	
HEWLETT PACKARD COMPANY P O BOX 272400, 3404 E. HARMONY ROAD INTELLECTUAL PROPERTY ADMINISTRATION FORT COLLINS, CO 80527-2400			GOLDEN, JAMES R	
			ART UNIT	PAPER NUMBER
			2187	

DATE MAILED: 01/04/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	10/698,265	KARLSSON ET AL.	
	Examiner	Art Unit	
	James Golden	2187	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) Responsive to communication(s) filed on 30 October 2003.

2a) This action is **FINAL**. 2b) This action is non-final.

3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) Claim(s) 1-39 is/are pending in the application.
4a) Of the above claim(s) _____ is/are withdrawn from consideration.
5) Claim(s) _____ is/are allowed.
6) Claim(s) 1-39 is/are rejected.
7) Claim(s) _____ is/are objected to.
8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) The specification is objected to by the Examiner.

10) The drawing(s) filed on 30 October 2003 is/are: a) accepted or b) objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
a) All b) Some * c) None of:
1. Certified copies of the priority documents have been received.
2. Certified copies of the priority documents have been received in Application No. _____.
3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) Notice of References Cited (PTO-892)
2) Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date 02/09/04.

4) Interview Summary (PTO-413)
Paper No(s)/Mail Date. ____ .
5) Notice of Informal Patent Application (PTO-152)
6) Other: ____ .

DETAILED ACTION

The instant application 10/698264 has a total of 39 claims pending. There are 6 independent claims and 33 dependent claims. Claims 1-39 have been rejected in view of prior art.

Information Disclosure Statement

1. The information disclosure statement submitted on 02/09/2004 is in compliance with the provisions of 37 CFR 1.97. Accordingly, the information disclosure statement is being considered by the examiner.

Drawings

2. The drawings are objected to as failing to comply with 37 CFR 1.84(p)(5) because they include the following reference character(s) not mentioned in the description: 712 of Fig. 7A.

3. Corrected drawing sheets in compliance with 37 CFR 1.121(d) are required in reply to the Office action to avoid abandonment of the application. Any amended replacement drawing sheet should include all of the figures appearing on the immediate prior version of the sheet, even if only one figure is being amended. The figure or figure number of an amended drawing should not be labeled as "amended." If a drawing figure is to be canceled, the appropriate figure must be removed from the replacement sheet, and where necessary, the remaining figures must be renumbered and appropriate changes made to the brief description of the several views of the drawings for consistency. Additional

replacement sheets may be necessary to show the renumbering of the remaining figures. Each drawing sheet submitted after the filing date of an application must be labeled in the top margin as either "Replacement Sheet" or "New Sheet" pursuant to 37 CFR 1.121(d). If the changes are not accepted by the examiner, the applicant will be notified and informed of any required corrective action in the next Office action. The objection to the drawings will not be held in abeyance.

Specification

4. The title of the invention is not descriptive. A new title is required that is clearly indicative of the invention to which the claims are directed. The title "Method of Determining Bounds for Minimum Cost and Near-Optimal Location of Replica Data in Network Nodes" is suggested.

5. The disclosure is objected to because of the following informalities: the Application Numbers of the related applications are not given (page 1, line 4) and should be listed as --10698182, 10698264 and 10698265--; "(filed on the same day as this application)" (page 1, lines 5-6) should be corrected to --10/30/2003--; "allowable lime" (page 6, line 9) should be corrected. Appropriate correction is required.

6. The examiner respectfully requests that applicant correct the mention of figures in the disclosure (figure 1, figure 2, etc.) to read --Fig. 1-- and --Fig. 2--.

Claim Objections

7. **Claim 29** is objected to because of the following informalities: "an upper limit data transmission rate" should be corrected to --an upper limit on data transmission rate--. Appropriate correction is required.

Claim Rejections - 35 USC § 101

8. 35 U.S.C. 101 reads as follows:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

9. **Claims 1-18** are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter. A method requires a practical application and a tangible result, neither of which are present in claims 1 or 2. Claims 3-18 are rejected because of their dependence upon claim 2. Claims 10-14 would be directed towards statutory subject matter if claim 10 rewritten in independent form including all of the limitations of the base claim and any intervening claims.

Claim Rejections - 35 USC § 102

10. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

11. **Claims 1-3, 4, 6-8, 9-16, 19-20, 22-23, 27, 31-32 and 37-39 are rejected**

under 35 U.S.C. 102(b) as being anticipated by Karlsson et al. ("Do We Need

Replica Placement Algorithms in Content Delivery Networks?").

12. **With respect to claim 1**, Karlsson et al. disclose a method of instantiating a data placement heuristic for a distributed storage system comprising the steps of:

- receiving heuristic parameters (page 3 [page 1 has abstract and introduction], column 1, paragraphs 2-7); and
- running an algorithm on a node of the distributed storage system which instantiates a particular data placement heuristic selected from a range of data placement heuristics according to the heuristic parameters (page 2, column 2, paragraph 4).

13. **With respect to claim 2**, Karlsson et al. disclose a method of instantiating a data placement heuristic for a distributed storage system comprising the steps of:

- receiving heuristic parameters (table 2; page 3, column 2, paragraph 4) which comprise a cost function (table 1; page 3, column 2, paragraph 3), a placement constraint (page 3, column 1, paragraph 8, lines 1-4), a metric scope (page 3, column 1, paragraph 9, line 4 -- column 2, paragraph 1), an approximation technique (page 3, column 2, paragraph 2), and an evaluation interval (page 2, column 1, paragraph 2); and

- running an algorithm on a node of the distributed storage system which instantiates a particular data placement heuristic selected from a range of data placement heuristics according to the heuristic parameters (page 3, column 1, paragraph 9, line 8 -- column 2, paragraph 1, line 2; "If... only one node is considered, the heuristic is decentralized and has to be run everywhere" entails that the heuristic is run on all nodes of the system).

14. **With respect to claim 3**, Karlsson et al. disclose the method of claim 2 (see above paragraph 13) wherein the heuristic parameters further comprise at least one additional placement constraint (page 3, column 2, paragraph 3, lines 4-6).

15. **With respect to claim 4**, Karlsson et al. disclose the method of claim 2 (see above paragraph 13) wherein the heuristic parameters further comprise a routing knowledge parameter (page 3, column 1, paragraph 7).

16. **With respect to claim 6**, Karlsson et al. disclose the method of claim 2 (see above paragraph 13) wherein the metric scope comprises a node scope, which ranges from a single node of the distributed storage system to all distributed storage system (page 3, column 1, paragraph 9, line 4 -- column 2, paragraph 1).

17. **With respect to claim 7**, Karlsson et al. disclose the method of claim 6 (see above paragraph 16) wherein the metric scope further comprises a client scope, which ranges from local clients accessing the single node to all clients accessing all of the nodes (page 3, column 1, paragraph 9, line 4 -- column 2, paragraph 1).

18. **With respect to claim 8**, Karlsson et al. disclose the method of claim 6 (see above paragraph 16) wherein the metric scope further comprises a data object scope, which ranges from local data objects stored on the single node to all data objects stored on all of the nodes (page 3, column 1, paragraph 9, line 4 -- column 2, paragraph 1).

19. **With respect to claim 9**, Karlsson et al. disclose the method of claim 2 (see above paragraph 13) wherein the heuristic parameters further comprise at least one additional approximation technique (page 3, column 2, paragraph 2).

20. **With respect to claim 10**, Karlsson et al. disclose the method of claim 2 (see above paragraph 13) wherein the algorithm comprises a determination of costs for nodes and data objects within the metric scope according to the cost function (page 3, column 2, paragraph 2, lines 4-8; page 6, column 1, paragraph 3).

21. **With respect to claim 11**, Karlsson et al. disclose the method of claim 10 (see above paragraph 20) wherein the algorithm further comprises a sort of the costs (page 3, column 2, paragraph 2, lines 4-8).

22. **With respect to claim 12**, Karlsson et al. disclose the method of claim 11 (see above paragraph 21) wherein the approximation comprises a ranking technique (page 3, column 2, paragraph 2, lines 4-8).

23. **With respect to claim 13**, Karlsson et al. disclose the method of claim 12 (see above paragraph 22) wherein the ranking technique comprises a greedy algorithm (page 3, column 2, paragraph 2, lines 4-8).

24. **With respect to claim 14**, Karlsson et al. disclose the method of claim 11 (see above paragraph 21) wherein the approximation technique comprises a threshold technique (page 3, column 2, paragraph 2, lines 4-8; the constraints represent thresholds).

25. **With respect to claim 15**, Karlsson et al. disclose the method of claim 2 (see above paragraph 13) wherein the approximation technique comprises a random technique (page 3, column 2, paragraph 2, lines 12-16; the improvement approximation stopped after the initial random seeding is a random technique).

26. **With respect to claim 16**, Karlsson et al. disclose the method of claim 2 (see above paragraph 13) wherein the approximation technique comprises an improvement technique (page 3, column 2, paragraph 2, lines 12-20).

27. **With respect to claim 19**, Karlsson et al disclose a method of instantiating a data placement heuristic for a distributed storage system comprising the steps of:

- receiving heuristic parameters (table 2; page 3, column 2, paragraph 4) which comprise a cost function (table 1; page 3, column 2, paragraph 3), a placement constraint (page 3, column 1, paragraph 8, lines 1-4), a metric scope (page 3, column 1, paragraph 9, line 4 -- column 2, paragraph 1), an approximation technique (page 3, column 2, paragraph 2), and an evaluation interval (page 2, column 1, paragraph 2); and
- for each evaluation interval within an execution of the data placement heuristic, performing the steps of:

- for data objects and nodes within the metric scope, determining a cost array which comprises cost elements determined from the cost function (page 3, column 2, paragraph 2, lines 4-8; all of the costs compose the array);
- if the approximation technique comprises a ranking technique, placing the data objects according to the placement constraint in order of rank within a sorted cost array (page 3, column 2, paragraph 2, lines 4-8);
- if the approximation technique comprises a greedy ranking technique, placing the data objects according to the placement constraint in order of the rank within the sorted cost array, and recomputing and resorting the cost array for remaining cost elements after each placement decision (page 3, column 2, paragraph 2, lines 11-21); and
- if the approximation technique comprises a threshold technique, eliminating the cost elements of the cost array which violate a threshold, thereby forming a threshold limited cost array, and placing the data objects according to the placement constraint in order of rank within a sorted version of the threshold limited cost array (page 3, column 2, paragraph 2, lines 4-8; the constraints represent thresholds, and the costs that are not ranked because they violate a constraint compose the threshold limited cost array).

28. **With respect to claim 20**, Karlsson et al. disclose the method of claim 19

(see above paragraph 27) wherein the cost function comprises a rate of read accesses by a client for a particular data object (page 3, column 1, paragraph 2; table 1; page 3, column 2, paragraph 3).

29. **With respect to claim 22**, Karlsson et al. disclose the method of claim 19

(see above paragraph 27) wherein the cost function comprises a measure of a distance between a particular client and a particular node (page 3, column 1, paragraph 3; table 1; page 3, column 2, paragraph 3).

30. **With respect to claim 23**, Karlsson et al. disclose the method of claim 19

(see above paragraph 27) wherein the cost function comprises a storage cost for storing a particular object on a particular node (page 3, column 1, paragraph 2, lines 4-6; “the cost of placing one extra object on one node” is interpreted as storage cost).

31. **With respect to claim 27**, Karlsson et al. disclose the method of claim 19

(see above paragraph 27) wherein the placement constraint comprises a storage capacity constraint, which places an upper limit on a storage capacity for a node (page 3, column 1, paragraph 8, lines 4-6).

32. **With respect to claim 31**, Karlsson et al. disclose the method of claim 19

(see above paragraph 27) wherein the placement constraint comprises replica constraint, which places an upper limit on a number of replicas that can be placed on nodes of the distributed storage system (page 3, column 1, paragraph 8, lines 1-4).

33. **With respect to claim 32**, Karlsson et al. disclose the method of claim 19 (see above paragraph 27) wherein the placement constraint comprises a delay constraint, which places an upper limit on a response time for requests within the distributed storage system (page 5, column 1, paragraph 2, line 8 -- column 2, paragraph 1; "a response latency below Y msec" is the threshold).

34. **With respect to claim 37**, Karlsson et al. disclose a computer readable memory comprising computer code for implementing a method of instantiating a data placement heuristic for a distributed storage system, the method of instantiating the data placement heuristic comprising the steps of:

- receiving heuristic parameters (page 3 [page 1 has abstract and introduction], column 1, paragraphs 2-7); and
- running an algorithm on a node of the distributed storage system which instantiates a particular data placement heuristic selected from a range of data placement heuristics according to the heuristic parameters (page 2, column 2, paragraph 4).

35. **With respect to claim 38**, Karlsson et al. disclose a computer readable memory comprising computer code for implementing a method of instantiating a data placement heuristic for a distributed storage system, the method of instantiating data placement heuristic comprising the steps of:

- receiving heuristic parameters (table 2; page 3, column 2, paragraph 4) which comprise a cost function (table 1; page 3, column 2, paragraph 3), a placement constraint (page 3, column 1, paragraph 8, lines 1-4), a metric scope (page 3, column 1, paragraph 9, line 4 -- column 2, paragraph 1),

an approximation technique (page 3, column 2, paragraph 2), and an evaluation interval (page 2, column 1, paragraph 2); and

- running an algorithm on a node of the distributed storage system which instantiates a particular data placement heuristic selected from a range of data placement heuristics according to the heuristic parameters (page 3, column 1, paragraph 9, line 8 -- column 2, paragraph 1, line 2; "If... only one node is considered, the heuristic is decentralized and has to be run everywhere" entails that the heuristic is run on all nodes of the system).

36. **With respect to claim 39**, Karlsson et al disclose a computer readable memory comprising computer code for implementing a method of instantiating a data placement heuristic for a distributed storage system, the method of instantiating the data placement heuristic comprising the steps of:

- receiving heuristic parameter (table 2; page 3, column 2, paragraph 4) which comprise a cost function (table 1; page 3, column 2, paragraph 3), a placement constraint (page 3, column 1, paragraph 8, lines 1-4), a metric scope (page 3, column 1, paragraph 9, line 4 -- column 2, paragraph 1), an approximation technique (page 3, column 2, paragraph 2), and an evaluation interval (page 2, column 1, paragraph 2); and
- for each evaluation interval within an execution of the data placement heuristic, performing the steps of:
 - for data objects and nodes within the metric scope, determining a cost array which comprises cost elements determined from the cost

function (page 3, column 2, paragraph 2, lines 4-8; all of the costs compose the array);

- if the approximation technique comprises a ranking technique, placing the data objects according to the placement constraint in order of rank within a sorted cost array (page 3, column 2, paragraph 2, lines 4-8);
- if the approximation technique comprises a greedy ranking technique, placing the data objects according to the placement constraint in order of the rank within the sorted cost array, and recomputing and resorting the cost array for remaining cost elements after each placement decision (page 3, column 2, paragraph 2, lines 11-21); and
- if the approximation technique comprises a threshold technique, eliminating the cost elements of the cost array which violate a threshold, thereby forming a threshold limited cost array, and placing the data objects according to the placement constraint in order of rank within a sorted version of the threshold limited cost array (page 3, column 2, paragraph 2, lines 4-8; the constraints represent thresholds, and the costs that are not ranked because they violate a constraint compose the threshold limited cost array).

37. **Claims 2-3, 17-19, 21, 24, 26, 28-30 and 32-36** are rejected under 35 U.S.C. 102(b) as being anticipated by Karlsson et al. ("A Framework for Evaluating Replica Placement Algorithms").

38. **With respect to claim 2**, Karlsson et al. disclose a method of instantiating a data placement heuristic for a distributed storage system comprising the steps of:

- receiving heuristic parameters (page 2, column 2, paragraph 3 -- page 3, column 1, paragraph 9) which comprise a cost function (table I; page 3, column 1, paragraph 10), a placement constraint (page 3, column 1, paragraph 6, "Number of Replicas"), a metric scope (page 5, column 1, paragraph 2), an approximation technique (page 5, column 1, paragraph 3 -- column 2, paragraph 5), and an evaluation interval (page 6, column 2, paragraph 2 -- page 7, column 1, paragraph 3); and
- running an algorithm on a node of the distributed storage system which instantiates a particular data placement heuristic selected from a range of data placement heuristics according to the heuristic parameters (page 2, column 1, paragraph 2, lines 8-11; caption for table III, "the node that makes the decision each time it runs the algorithm").

39. **With respect to claim 3**, Karlsson et al. disclose the method of claim 2 (see above paragraph 38) wherein the heuristic parameters further comprise at least one additional placement constraint (.

40. **With respect to claim 17**, Karlsson et al. disclose the method of claim 2 (see above paragraph 38) wherein the approximation technique comprises a hierarchical technique (page 5, column 2, paragraph 4).

41. **With respect to claim 18**, Karlsson et al. disclose the method of claim 2 (see above paragraph 38) wherein the approximation technique comprises a multiphase technique (page 5, column 2, paragraph 5).

42. **With respect to claim 19**, Karlsson et al. disclose a method of instantiating a data placement heuristic for a distributed storage system comprising the steps of:

- receiving heuristic parameters (page 2, column 2, paragraph 3 -- page 3, column 1, paragraph 9) which comprise a cost function (table I; page 3, column 1, paragraph 10), a placement constraint (page 3, column 1, paragraph 6, "Number of Replicas"), a metric scope (page 5, column 1, paragraph 2), an approximation technique (page 5, column 1, paragraph 3 -- column 2, paragraph 5), and an evaluation interval (page 6, column 2, paragraph 2 -- page 7, column 1, paragraph 3); and
- for each evaluation interval within an execution of the data placement heuristic, performing the steps of:
 - for data objects and nodes within the metric scope, determining a cost array which comprises cost elements determined from the cost function (table I; page 3, column 1, paragraph 10; page 5, column 1, paragraph 2);
 - if the approximation technique comprises a ranking technique, placing the data objects according to the placement constraint in order of rank within a sorted cost array (page 5, column 1, paragraph 4);

- if the approximation technique comprises a greedy ranking technique, placing the data objects according to the placement constraint in order of the rank within the sorted cost array, and recomputing and resorting the cost array for remaining cost elements after each placement decision (page 5, column 1, paragraph 4; resorting is inherent in ranking part of greedy algorithm); and
- if the approximation technique comprises a threshold technique, eliminating the cost elements of the cost array which violate a threshold, thereby forming a threshold limited cost array, and placing the data objects according to the placement constraint in order of rank within a sorted version of the threshold limited cost array (page 5, column 1, paragraph 5; page 5, column 2, paragraph 7; table II; table II shows that two threshold heuristics, Awerbuch FAP and Dist Awerbuch, use the R(plain) approximation method in addition to a threshold, and R(plain) is the approximation method used for all of the ranking algorithms, so these threshold algorithms rank the cost array as well).

43. **With respect to claim 21**, Karlsson et al. disclose the method of claim 19 (see above paragraph 42) wherein the cost function comprises a rate of write accesses by a client for a particular object (page 2, column 2, paragraph 5; page 3, column 1, paragraph 10; table I, group 3).

44. **With respect to claim 24**, Karlsson et al. disclose the method of claim 19 (see above paragraph 42) wherein the cost function comprises an access time comprising a timestamp indicating a most recent time when a particular data object was accessed at a particular node (page 2, column 2, paragraph 9).

45. **With respect to claim 26**, Karlsson et al. disclose the method of claim 19 (see above paragraph 42) wherein the cost function comprises a hit ratio of transparent en route caches along a path from a client to a node (page 2, column 2, paragraph 10).

46. **With respect to claim 28**, Karlsson et al. disclose the method of claim 19 (see above paragraph 42) wherein the placement constraint comprises a load capacity constraint, which places an upper limit on a rate of requests for a node (page 3, column 1, paragraph 3).

47. **With respect to claim 29**, Karlsson et al. disclose the method of claim 19 (see above paragraph 42) wherein the placement constraint comprises a node bandwidth capacity, which places an upper limit data transmission rate for a node (page 3, column 1, paragraph 4).

48. **With respect to claim 30**, Karlsson et al. disclose the method of claim 19 (see above paragraph 42) wherein the placement constraint comprises an upper limit on bandwidth between first and second nodes (page 3, column 1, paragraph 5).

49. **With respect to claim 32**, Karlsson et al. disclose the method of claim 19 (see above paragraph 42) wherein the placement constraint comprises a delay

constraint, which places an upper limit on a response time for requests within the distributed storage system (page 3, column 1, paragraph 8).

50. **With respect to claim 33**, Karlsson et al. disclose the method of claim 19 (see above paragraph 42) wherein the placement constraint comprises an availability constraint, which places a lower limit on availability of data objects within the distributed storage system (page 3, column 1, paragraph 9).

51. **With respect to claim 34**, Karlsson et al. disclose the method of claim 19 (see above paragraph 42) wherein the approximation technique further comprises a hierarchical technique (page 5, column 2, paragraph 4).

52. **With respect to claim 35**, Karlsson et al. disclose the method of claim 19 (see above paragraph 42) wherein the approximation technique further comprises a multiphase technique (page 5, column 2, paragraph 5).

53. **With respect to claim 36**, Karlsson et al. disclose the method of claim 35 (see above paragraph 52) wherein the multiphase technique comprises an improvement technique (page 5, column 2, paragraph 5).

Claim Rejections - 35 USC § 103

54. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

55. **Claim 5** is rejected under 35 U.S.C. 103(a) as being unpatentable over Karlsson et al. ("Do We Need Replica Placement Algorithms in Content Delivery Networks?") in view of Dahlin et al. (US 2004/0064577).

56. **With respect to claim 5**, Karlsson et al. teach the method of claim 2 (see above paragraph 13). Karlsson et al. fail to disclose the limitation wherein the heuristic parameters further comprise an activity history parameter.

However, Dahlin et al. disclose the limitation wherein the heuristic parameters further comprise an activity history parameter [022, lines 11-21; 194, lines 9-16; 256; the system uses heuristics for data placement, and keeps a history of accesses to create a hint, or prefetch, list of data objects].

Karlsson et al. and Dahlin et al. are analogous art because they are from the same field of endeavor, namely the storage of replicated data.

At the time of invention it would have been obvious to a person of ordinary skill in the art to implement the history parameter of Dahlin et al. with the replication placement algorithms of Karlsson et al. The motivation for doing so would have been to "to determine an estimate of probability of one or more data objects on at least one server being requested as a demand request" [022, lines 8-11].

Therefore, it would have been obvious to combine Dahlin et al. with Karlsson et al. for the benefit of replication placement algorithms that created prefetch lists to obtain the invention as specified in claim 5.

57. **Claim 25** is rejected under 35 U.S.C. 103(a) as being unpatentable over Karlsson et al. ("Do We Need Replica Placement Algorithms in Content Delivery Networks?") in view of Cousins et al. ("The Embedded Genetic Allocator").

58. **With respect to claim 25**, Karlsson et al. disclose the method of claim 19 (see above paragraph 27). Karlsson et al. do not disclose the limitation wherein the cost function comprises a load time comprising a timestamp indicating a time when a particular data object was stored on a particular node.

However, Cousins et al. disclose the limitation wherein the cost function comprises a load time comprising a timestamp indicating a time when a particular data object was stored on a particular node (page 2168, paragraphs 2-4; although only a "process" timestamp is disclosed, this could include a data storage operation).

Karlsson et al. and Cousins et al. are analogous art because they are from the same field of endeavor, namely the storage of data on memory nodes.

At the time of invention it would have been obvious to a person of ordinary skill in the art to include the load timestamp of Cousins et al. in the cost function of Karlsson et al. The motivation for doing so would have been because "the GA [genetic algorithm] engine uses the event log to calculated the durations of the states in all the traces. It then calculates the overall fitness score of the run by combining the state timing information according to specified user criteria" (page 2168, paragraph 4, lines 1-4)

Therefore, it would have been obvious to combine Cousins et al. with Karlsson et al. for the benefit of a cost function that takes into account the load timestamp to obtain the invention as in claim 25.

Conclusion

59. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

- Borowsky et al. (US 6,321,317) teaches the use of heuristics in organizing a data storage network.
- Lumelsky et al. (US 6,463,454 and US 6,466,980) teaches a system in which heuristics are used to manage replicas.
- Richardson et al. (US 6,249,802) teaches a method for distributing replicas in a data storage network.
- Rabinovich (6,256,675) teaches using heuristics for replica placement in a network.
- Karlsson et al. (US 2004/0034744) teaches a data storage system using caching instead of replica placement.
- Cousins, Loomis, Roeber, Schoeppner, Tobin, "The Embedded Genetic Allocator," 1998, IEEE International Conference on Systems, Man and Cybernetics, pages 2166-2171.
- Karlsson et al., "Choosing Replica Placement Heuristics for Wide-Area Systems," 2004, IEEE Conference on Distributed Computing Systems, pages 350-359.

- Bartolini et al., "Optimal Replica Placement in Content Delivery Networks," September 28, 2003, IEEE Conference on Networks, pages 125-130.

60. Any inquiry concerning this communication or earlier communications from the examiner should be directed to James Golden whose telephone number is 571-272-5628. The examiner can normally be reached on Monday-Friday, 8:30 AM - 5:30 PM.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Donald Sparks can be reached on 571-272-4201. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

James R. Golden
Patent Examiner
Art Unit 2187

MJ

B. R. Peal
Brian R. Peal
Primary Examiner
AU 2187
Reb 26, 2005